

AI/ML Challenge Tech Talk

Ryan Day

December 2020

CHIEF TECHNOLOGY OFFICER

Today I'll be sharing some of the technical and data science information that we learned from our recent AI/ML challenge.

The OCTO works to promote emerging technology and spread understanding of technology in the agency.

Our Challenge Event

challenge.gov/challenge/GSA-artificial-intelligence-AI-machine-learning-ML-challenge/



General Services
Administration

Page Contents

Description

Prizes

Rules

Judging Criteria

How To Enter

GSA Artificial Intelligence and Machine Learning End-User License Agreement Challenge 2020

Create an artificial intelligence or machine learning solution to aid in the review of end-user license agreements.

CHALLENGE DETAILS

-  TOTAL CASH PRIZES OFFERED: \$20,000
-  TYPE OF CHALLENGE: Software and apps
-  SUBMISSION START: 07/06/2020 10:00 AM ET
-  SUBMISSION END: 08/20/2020 5:00 PM ET

"Emerging technologies such as artificial intelligence and machine learning are paving the way to automate business processes ... Let us ride the digital wave of transformation together." – Keith Nakasone, Deputy Assistant Commissioner for Acquisition, General Services Administration (GSA)

"AI & ML leveraged successfully are force multipliers; transitioning historically resource intensive activities into high-speed-low-drag transactions. Huge thank you in advance for your participation and contributions!" – Sean C. Zerges, Director, Office of the Chief Technology Officer, GSA

CHIEF TECHNOLOGY OFFICER

During the summer of 2020, the OCTO and FAS ITC collaborated to host an online machine learning challenge.

Full details are at this website:

<https://www.challenge.gov/challenge/GSA-artificial-intelligence-AI-machine-learning-ML-challenge/>

What is an End User License Agreement (EULA)?

COMPANY End User License Master Agreement

COMPANY END USER LICENSE MASTER AGREEMENT

This COMPANY Inc internet service agreement for software license and services ("**COMPANY End User License Master Agreement**") is made as of the Effective Date between COMPANY, Inc., a LOCATION corporation located at ADDRESS ("COMPANY"), and the party/signatory to the attached Purchase Order ("**Customer**"). For good and valuable consideration, the parties hereto, intending to be legally bound, hereby agree as follows:

1. DEFINITIONS.

1.1 "Customer" means the persons, entity or agents and authorized representatives accepting this agreement.

2. OWNERSHIP.

All Content including any and all intellectual property rights in the Content are owned by COMPANY, and Customer shall make no claim of ownership to any content, including subsequent versions or enhancements to Content made at Customer's request that are implemented by COMPANY or its licensors. This Agreement does not constitute a Copyright license. COMPANY warrants that is the lawful

CHIEF TECHNOLOGY OFFICER

Let's start with the business problem wanted to solve:

The goal of this challenge is to develop an artificial intelligence (AI), or machine learning (ML) solution that will review end-user license agreements (EULA) for terms and conditions that are unacceptable to the government. On average it takes all parties involved approximately 7-14 days to review an EULA and ensure that all unacceptable terms and conditions are identified.

A EULA details the rights and restrictions which apply to the use of software or services. It can also be known as a software license agreement or acceptable use policy. As part of the acquisition process of software or services, GSA reviews the associated EULAs. This review must be completed prior to awarding new contracts or modifying existing contracts. A GSA contracting officer (CO) reviews applicable EULAs for terms and conditions that are not in accordance with Federal law and regulations. The CO may also coordinate a legal review with the Office of General Counsel if they feel it is warranted. Should EULAs contain language that would conflict with Federal law and regulations, the CO must negotiate changes to the EULA to remove the problematic language.

We are looking for a solution that will use AI and/or ML to improve this

manual process. The solution will include a user interface that GSA will use to process the documents and identify unacceptable clauses/terms in the EULAs. [Watch our AI / ML challenge video](#) to learn more about the desired solution.

This solution will decrease the time spent manually reviewing EULAs and free resources to focus on other aspects of the acquisition process. It will also improve the accuracy and consistency of the review process.

Why would clauses be unacceptable to government?

Attachment B

Schedule 70 EULA Matrix

Both GSA and Government ordering activities placing orders under GSA Schedule 70 contracts are required to comply with the FAR clause at 12.212(a), which provides, in relevant part, that commercial computer software and documentation shall be acquired under licenses customarily provided to the public "to the extent such licenses are consistent with Federal law and otherwise satisfy the Government's needs."

Below is a list of terms and conditions commonly occurring in software manufacturers' unmodified commercial agreements that are inconsistent with Federal law and Government needs and, therefore, with FAR 12.212(a). The terms and conditions listed below are non-compliant regardless of the type of agreement: end-user license agreements (EULAs), maintenance agreements, terms of service (TOS), etc.

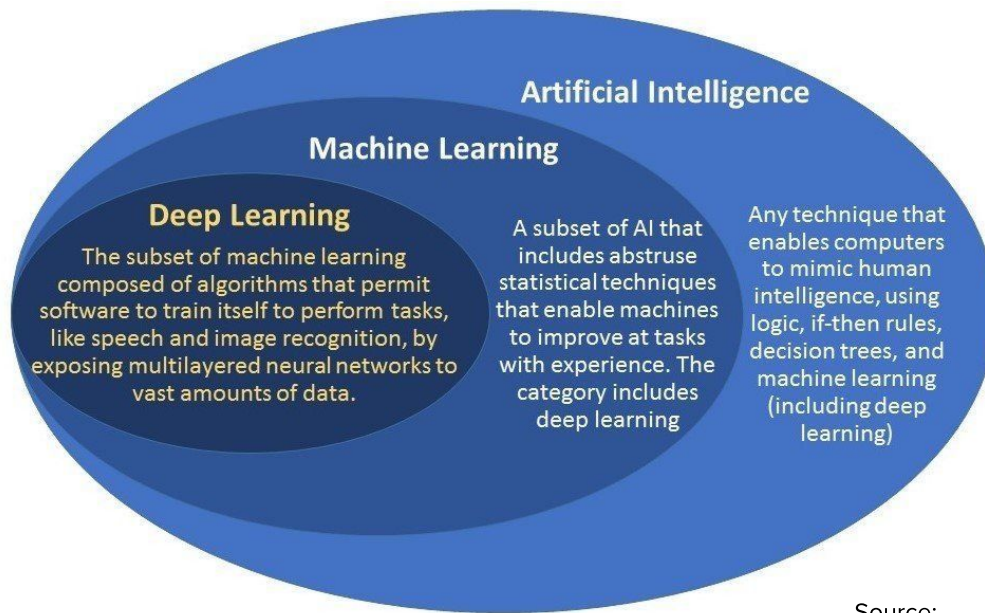
In order to avoid delays caused by legal review and negotiation of each individual set of terms, or loss of business that occurs when the terms are added on Schedule unmodified and ordering activities later decline to place an order because of non-conforming terms, manufacturers (and/or dealers or resellers, where appropriate) should create compliant agreements that do not contain the clauses listed below.

	Terms and conditions	Problem/recommendation
1	Definition of contracting parties	The Government customer (licensee), under GSA Schedule contracts, is the "ordering activity," defined as an "entity authorized to order under GSA Schedule contracts as defined in GSA Order ADM4800.2G, as may be revised from time to time." The licensee or customer cannot be an individual because any implication of individual licensing triggers the requirement for legal review by Federal employee unions. Conversely, because of competition rules, the contractor must be defined as a single entity even if the contractor is part of a corporate group. The Government cannot contract with the group, or in the alternative with a set of contracting parties.
2	Contract formation via usage.	Under FAR 1.601(a), in an acquisition involving the use of

CHIEF TECHNOLOGY OFFICER

There are a variety of reasons a clause might be unacceptable to government. We provided "Attachment B" to the teams in the reference materials:

<https://github.com/GSA/ai-ml-challenge-2020/tree/master/reference>




Source:
Data Science Dojo

CHIEF TECHNOLOGY OFFICER

This shows how AI & ML relate to each other.

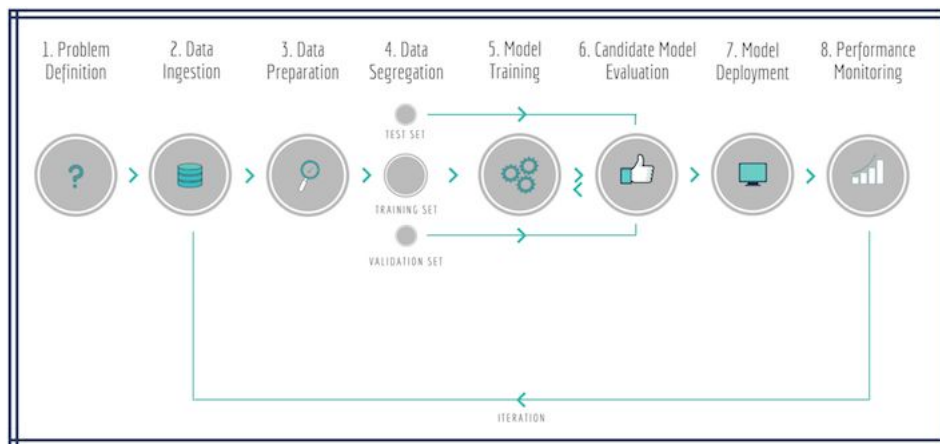
Machine Learning Flow



CHIEF TECHNOLOGY OFFICER

This is a typical machine learning flow.

Machine Learning Pipeline



Source: Semi Koen, Towards Data Science:

<https://towardsdatascience.com/architecting-a-machine-learning-pipeline-a847f094d1c7>

CHIEF TECHNOLOGY OFFICER


This is a view of a machine learning pipeline.

<https://towardsdatascience.com/architecting-a-machine-learning-pipeline-a847f094d1c7>

Our Challenge

Challenge.gov

Github.com



GSA Artificial Intelligence and Machine Learning End-User License Agreement Challenge 2020

Create an artificial intelligence or machine learning solution to aid in the review of end-user license agreements.

CHALLENGE DETAILS

- TOTAL CASH PRIZES OFFERED: \$20,000
- TYPE OF CHALLENGE: Software and apps
- SUBMISSION START: 07/06/2020 10:00 AM ET
- SUBMISSION END: 08/20/2020 5:00 PM ET

"Emerging technologies such as artificial intelligence and machine learning are paving the way to automate business processes ... Let us ride the digital wave of transformation together." – Keith Nakasone, Deputy Assistant Commissioner for Acquisition, General Services Administration (GSA)

"AI & ML leveraged successfully are force multipliers; transitioning historically resource intensive activities into high-speed-low-drag transactions. Huge thank you in advance for your participation and contributions!" – Sean C. Zerges, Director, Office of the Chief Technology Officer, GSA

AI/ML Challenge 2020

Registering for the Challenge

The first task is to register your team for the challenge. To do that, follow the instructions on the [challenge website](#).

Example EULA Documents

EULA ("yoola") documents are typically received as Microsoft Word or PDF documents. Examples of these are available in the [Reference](#) section.

Training Data For The Challenge

We are providing the initial set of training data that has been extracted from actual EULA documents and had identifying information, such as company name, removed. We may provide updates to the training data during the challenge, and list them in the table below.

Data set	Description	Date published
Training Data Set 1 v1	Initial batch of data. Known issues: clause text contains control characters, such as embedded Line Feed (\n) characters.	7/6/2020

Format of training data:

No packages published
[Publish your first package](#)

Contributors


- Ryandaydev Ryan Day
- tslu00 Tiffany Liu

We had instructions on challenge.gov and github.

Here were the github instructions:

<https://github.com/GSA/ai-ml-challenge-2020/blob/master/README.md>

Judging Criteria



CHIEF TECHNOLOGY OFFICER

These were the scoring criteria. We published a scoring rubric to explain them.

https://github.com/GSA/ai-ml-challenge-2020/blob/master/reference/AI_ML%20Challenge%20Scoring%20Rubric.pdf

Demo of Winning Solutions

#1 - Dev Tech

**#2 - Gaussian
Solutions**

#3 - Team SoKat

CHIEF TECHNOLOGY OFFICER

These were the three winners.



EULA Clause Analysis

Welcome to EULACheck, Dev Technology's End User License Agreement (EULA) analysis app!

You may upload a EULA in Word (.docx) or PDF (.pdf) format, or paste EULA content directly into the text box from your clipboard.

We will take a look at the **first 10 clauses**, and determine if the Federal Government is likely to object to the language there.

We will also show you clauses that closely match yours which are known to be unacceptable and/or acceptable to the Federal Government.

All content on this site is informational - it should not be construed as legal advice.

Finally, you can provide your feedback by telling us if our analysis was correct or not.

***The application works best on Chrome or Firefox**

Choose File | No file chosen

Submit EULA File

OR

" the Customer does not make any admissions (save where required by court order or governmental regulations, and where the Customer is required under the terms of such order or regulations not to first consult with the Company) which may be prejudicial to the defense or settlement of any Claim without the Company's approval (not to be unreasonably withheld or delayed)."

CHIEF TECHNOLOGY OFFICER

Requests. Company will notify Customer before Customer exceeds the Tile Request Use Limit indicated on the Order Form. If Customer exceeds its Tile Request Use Limits during the License Term, Company will invoice Customer for Overages on written notice (which may be by email). If, after 30 days from the date of that written notice, Customer continues to exceed its Tile Request Use Limit, Company may stop providing the Service to the Customerinitiate a claim with the Contracting Officer under the Contract Disputes Act.

Submit Clause Text

Clause	Acceptable?	Feedback	Closest Acceptable Match	Closest Unacceptable Match
Requests. Company will notify Customer before Customer exceeds the Tile Request Use Limit indicated on the Order Form. If Customer exceeds its Tile Request Use Limits during the License Term, Company will invoice Customer for Overages on written notice (which may be by email). If, after 30 days from the date of that written notice, Customer continues to exceed its Tile Request Use Limit, Company may stop providing the Service to the Customerinitiate a claim with the Contracting Officer under the Contract Disputes Act.	Acceptable <div>99%</div>	<div><input type="radio"/> Acceptable</div> <div><input type="radio"/> Unacceptable</div> <div><input type="radio"/> Not Sure</div> <div><input type="radio"/> Not a Clause</div>	Term of this Addendum. This Addendum will commence on the Addendum Effective Date and continue for a period of twelve months ("Initial Addendum Term"). Upon the effective date of termination of this Addendum in accordance with the Contract Disputes Act, Client's access to the Hosted Service provided pursuant to this Addendum (and all licenses granted under this Addendum) will cease and COMPANY will delete all backed-up Client Data from the Hosting Infrastructure within 30 days of termination of this Agreement.	Company warrants that the Service will, for a period of sixty (60) days from the date of your receipt, perform substantially in accordance with Service written materials accompanying it



Gaussian Solutions
Delivering Comprehensive System Solutions

GSA Challenge

Upload the document you want analyzed

Select file...

Browse

Analyze



CHIEF TECHNOLOGY OFFICER

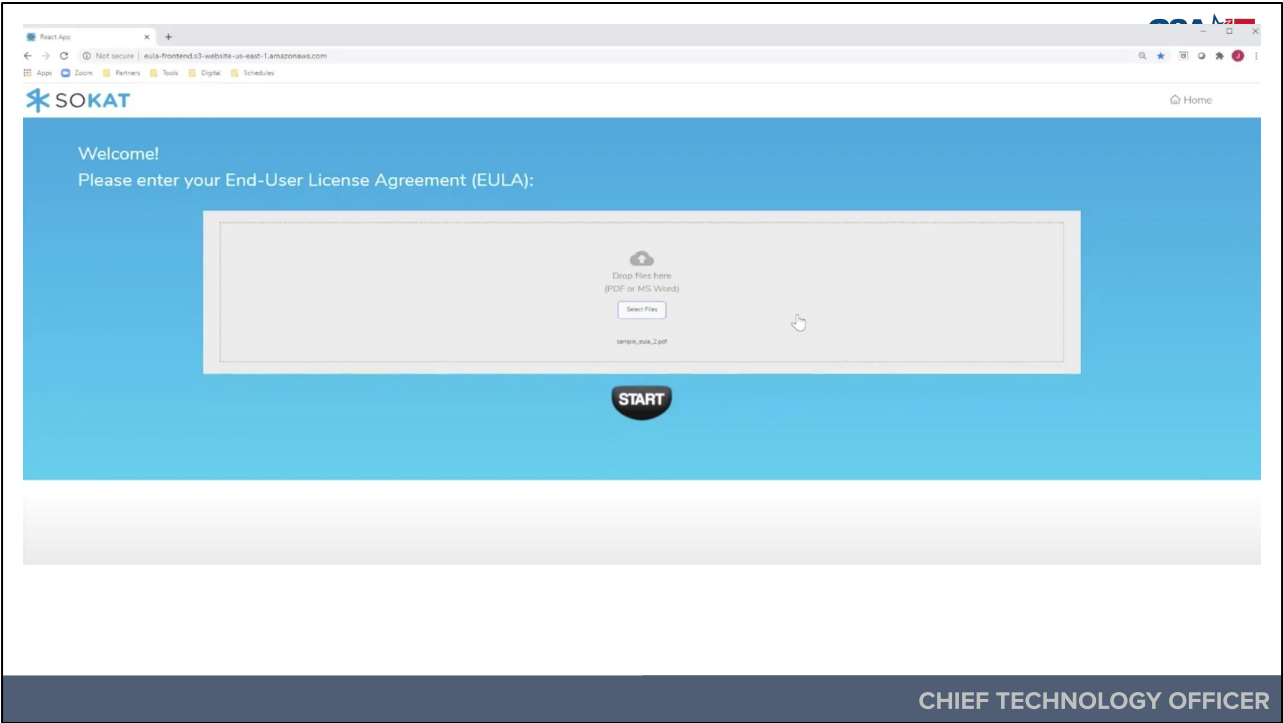
Showing analysis for sample_eula_1.docx

Analyze another document

Submit Reviewer's Decisions

Clause	Recommended	Confidence	Reviewer's Decision
MASTER SERVICES SUBSCRIPTION AGREEMENT	Accept	94.32	<input checked="" type="radio"/> Agree <input type="radio"/> Disagree
This Master Services Subscription Agreement (the "Agreement") sets forth the terms and conditions governing COMPANY' provision to Client of a cloud-based asset management and decision support system.	Accept	85.5	<input checked="" type="radio"/> Agree <input type="radio"/> Disagree
This Agreement, including the Order Form attached to it, as well as any Order Forms and Statements of Work entered into by the parties from time to time, the underlying AGENCY Schedule Contract, and Schedule Pricelist, together constitute the entire agreement of the parties and supersede any prior and contemporaneous oral or written understanding as to the parties' relationship and the subject matter hereof. In the event of any conflict or contradiction among the foregoing documents, the documents will control in the order listed in Contract Clause 552.212-4(s). This Agreement may only be amended in a writing signed by both parties.	Reject	39.03	<input checked="" type="radio"/> Agree <input type="radio"/> Disagree
This Agreement may be executed in two or more counterparts, each of which will be deemed an original for all purposes, and together will constitute one and the same document. Once signed, both parties agree that any	Accept	89.71	<input checked="" type="radio"/> Agree <input type="radio"/> Disagree

CHIEF TECHNOLOGY OFFICER



Your EULA results for sample_eula_2.pdf:

↓ Save

✉ Email Results

Submit User Decision

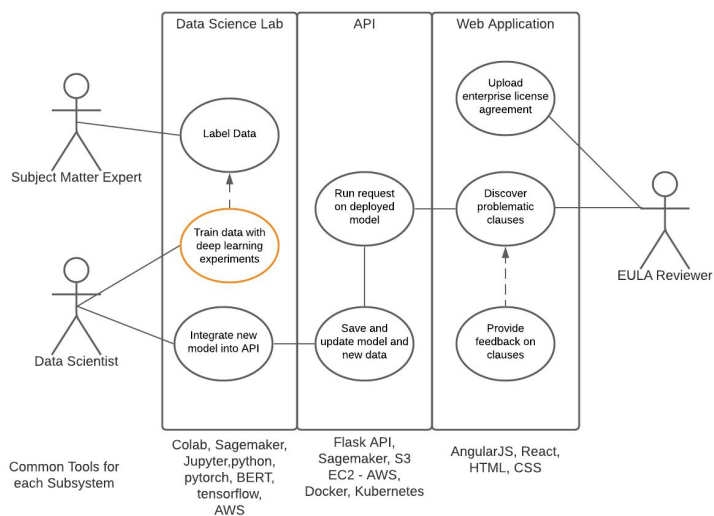
Clause Text	RF Score (%)	GRU Score (%)	XLNet Score (%)	Acceptability (Avg Score)	User Decision
1. DEFINITIONS	100	99	100	Yes (100%)	<input checked="" type="checkbox"/> Agree <input type="checkbox"/> Disagree
1.1 "Customer" means the persons, entity or agents and authorized representatives accepting this agreement.	92	100	100	Yes (97%)	<input checked="" type="checkbox"/> Agree <input type="checkbox"/> Disagree
1.2 "SOFTWARE" All content resulting and any all intellectual property rights in the Content are owned by COMPANY and Customer shall make no claim of ownership to any content including subsequent versions or enhancements to Content except as Customer's request that are implemented by COMPANY or its licensors. This Agreement does not constitute a Copyright license. COMPANY warrants that it is the lawful owner or licensee of all Content made accessible to Customer under this Agreement.	97	100	100	Yes (99%)	<input checked="" type="checkbox"/> Agree <input type="checkbox"/> Disagree
2. CONTENT CONTROL	100	100	100	Yes (100%)	<input checked="" type="checkbox"/> Agree <input type="checkbox"/> Disagree
2.1 - Lawful Purpose. Customer will only use COMPANY's Product & services for lawful purposes	98	100	100	Yes (99%)	<input checked="" type="checkbox"/> Agree <input type="checkbox"/> Disagree
2.2 - The purpose of this software to remove sensitive data permanently from user device as per user preference. The device shall include Windows Desktop, Laptop, Tablet, Server and Android Smartphone and Tablets. It is user choice to go with the software or as to remove the data security from the device. This way user can protect the data from any malware. COMPANY warrants data with built in secure algorithm to remove data securely from user data and securely remove from the disk. In data user do not have permission	92	99	100	Yes (97%)	<input checked="" type="checkbox"/> Agree <input type="checkbox"/> Disagree
2.3 - This SOFTWARE PRODUCT is protected by Copyright & Patent laws and international copyright treaties, as well as other intellectual property laws and treaties. The SOFTWARE PRODUCT is licensed to user and not sold.	99	100	100	Yes (100%)	<input checked="" type="checkbox"/> Agree <input type="checkbox"/> Disagree
2.4 - COMPANY or its related entities do not warrant, make any title and copyrights in and to the SOFTWARE PRODUCT, including but not limited to any images, writings, text, animations, video, audio, music, text incorporated into the SOFTWARE PRODUCT, the accompanying printed materials, and any copies of the SOFTWARE PRODUCT are owned by COMPANY or its. The SOFTWARE PRODUCT is Client initial. CONFIDENTIAL, Page 1 of 8 COMPANY initial protected Client initial. CONFIDENTIAL, Page 2 of 8 COMPANY initial.	98	97	100	Yes (98%)	<input checked="" type="checkbox"/> Agree <input type="checkbox"/> Disagree
4. PAYMENTS	100	99	100	Yes (100%)	<input checked="" type="checkbox"/> Agree <input type="checkbox"/> Disagree
4.1 - Licensing/Service Fee. License Fee are due at the time of signing the agreement.	94	74	100	Yes (89%)	<input checked="" type="checkbox"/> Agree <input type="checkbox"/> Disagree
4.2 - Support and Maintenance Payment. Support and Maintenance is included as part of licensing fee during subscription validity period.	97	97	100	Yes (98%)	<input checked="" type="checkbox"/> Agree <input type="checkbox"/> Disagree
4.3 - Overdue Payments. Overdue payments of License Fee required by this Agreement shall accrue interest at the lower of law and one (1%) per month (1.5% per month) the				Yes (76%)	

CHIEF TECHNOLOGY OFFICER

[illegible]

How Do These Apps Work?

Data Science Deployment Architecture



CHIEF TECHNOLOGY OFFICER

We found a lot of similarities between the solutions we received.

This is a big-picture view of the typical architecture that teams used.

How Did The Teams Work?

 [Save](#)[✉ Email Results](#)

Submit User Decision

Clause Text	RF Score (%)	GRU Score (%)	XLNet Score (%)	Acceptability (Avg Score)	User Decision
1. DEFINITIONS	100	99	100	Yes (100%)	<div><div></div> Agree <div></div> Disagree</div>
11. Customer's warrants the person, entity or agents and authorized representatives accepting this agreement.	92	100	100	Yes (97%)	<div><div></div> Agree <div></div> Disagree</div>
12. CONFIDENTIALITY AND NON-DISCLOSURE. All information, including but not limited to, confidential information, trade secrets, and other confidential information, disclosed by Customer to Vendor, or by Vendor to Customer, shall be held in confidence by the receiving party and shall not be disclosed to any third party without the prior written consent of the disclosing party. This obligation shall survive the termination or expiration of this Agreement.	97	100	100	Yes (99%)	<div><div></div> Agree <div></div> Disagree</div>
13. CONFIRMATION OF TERMS	100	100	100	Yes (100%)	<div><div></div> Agree <div></div> Disagree</div>
13.1. License Purpose. Customer will only use COMPANY's Product & services for lawful purposes.	98	100	100	Yes (99%)	<div><div></div> Agree <div></div> Disagree</div>
13.2. The purpose of the software or service under license is to provide a tool for the user to manage the data and information of the user. The user will not use the software or service for any other purpose. The user will not use the software or service to create or store any data or information that is not the user's own. The user will not use the software or service to create or store any data or information that is not the user's own.	92	99	100	Yes (94%)	<div><div></div> Agree <div></div> Disagree</div>
13.3. The SOFTWARE PRODUCT is protected by Copyright & Patent and other intellectual property rights, and all other rights in the software and the product. The SOFTWARE PRODUCT is intended for use by the user and not for resale.	99	100	100	Yes (100%)	<div><div></div> Agree <div></div> Disagree</div>
13.4. Customer or its authorized representative is not permitted to use the SOFTWARE PRODUCT for any purpose other than the purpose for which it was developed, and to use the SOFTWARE PRODUCT for any purpose other than the purpose for which it was developed, and to use the SOFTWARE PRODUCT for any purpose other than the purpose for which it was developed.	96	97	100	Yes (98%)	<div><div></div> Agree <div></div> Disagree</div>
4. WARRANTIES	100	99	100	Yes (100%)	<div><div></div> Agree <div></div> Disagree</div>
4.1. Limitation of Remedies. Vendor's sole obligation is to provide the SOFTWARE PRODUCT as described in the SOFTWARE PRODUCT, and to provide the SOFTWARE PRODUCT as described in the SOFTWARE PRODUCT.	94	74	100	Yes (89%)	<div><div></div> Agree <div></div> Disagree</div>
4.2. Support and Maintenance Period. Support and Maintenance is provided as part of the license fee during the support and maintenance period.	97	99	100	Yes (98%)	<div><div></div> Agree <div></div> Disagree</div>
4.3. General Cancellation. Customer may cancel this Agreement at any time without any penalty or obligation to pay any fee or charge.	93	97	100	Yes (96%)	<div><div></div> Agree <div></div> Disagree</div>

CHIEF TECHNOLOGY OFFICER

I am going to demonstrate the process followed by one of the teams.

CHIEF TECHNOLOGY OFFICER

The team used Jupyter Notebook and Python for their data science work.

Finished Application Architecture

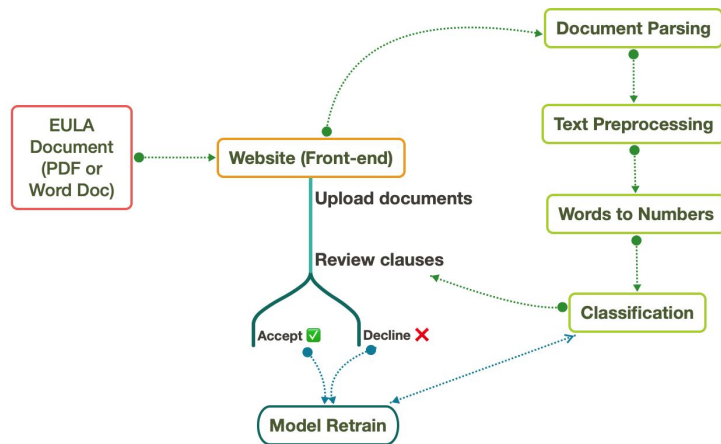
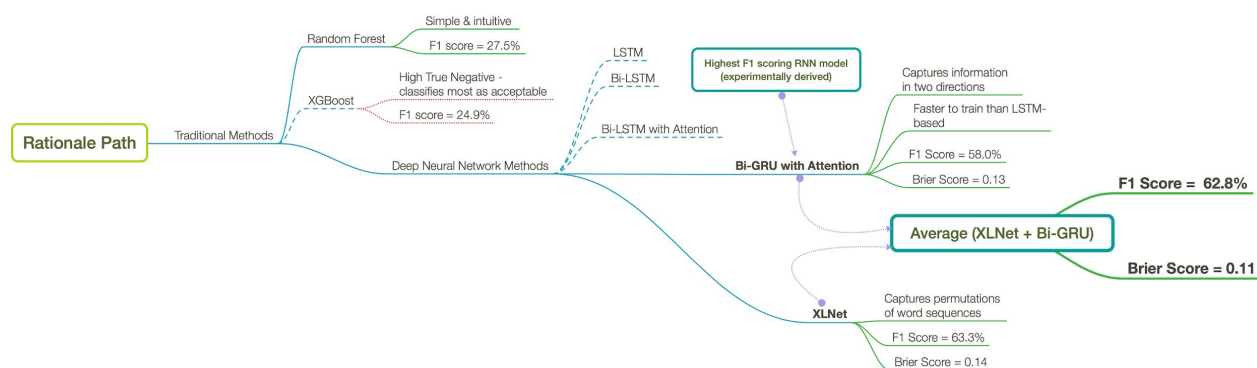


Figure 1: Overview of our solution presented as a flowchart.

CHIEF TECHNOLOGY OFFICER

This shows their finished architecture.

Rationale Map



CHIEF TECHNOLOGY OFFICER

This is how the team described their process for deciding on the machine learning libraries to use:

- As a first step, our team used traditional algorithms to approach the problem.
- We started with **Random Forest**, which is a popular machine learning algorithm that is widely used in classification tasks.
- **Random Forest** is simple and intuitive in nature. It does not require hyperparameter tuning and usually does not overfit to the dataset with an increase in the number of decision trees within the model. We achieved an F1 score of 27.5% with it.
- As the next step, we implemented the **XGBoost** algorithm. XGBoost is a tree-based algorithm (like Random Forest) but uses the technique of boosting. Boosting is an error-correction algorithm which gives a higher emphasis on data-points which are misclassified. Unlike random forests, the decision trees are created iteratively, where at each step, the tree puts more emphasis on the misclassified points, so as to reduce the overall error. We therefore used XGBoost as the natural next step to Random Forest. On testing however, the F1 score achieved through XGBoost was lower than Random Forest. This was potentially due to a high False Negative, resulting in a low Recall value. This means that the model was classifying most clauses

- (even the ones that were labelled unacceptable) as acceptable.
- After running several experiments trying to improve the accuracy of traditional models, we realized that more advanced, deep learning based models could potentially help us increase accuracies. We therefore started with the simplest form of sequence models that are used on textual data: **Recurrent Neural Networks (RNN)**. As a starting point we used a **Long Short-Term Memory (LSTM)** which attempts to resolve the vanishing gradient, a known obstacle for RNNs. LSTMs however capture the flow of information in one direction (left to right in case of sentences). **Bidirectional LSTMs** capture the flow of information from both left to right and right to left. This serves as an advantage as the model can learn from the future and the past information at a given point in the sentence. To add interpretability to our results, we added an **attention** layer that provides the importance of words in the decision making process of the model. The attention layer weighs those words differently providing more emphasis on words that have stronger relationships. Experiments run by researchers and practitioners have shown that **Gated Recurrent Units (GRUs)** train faster and provide a proportional (or sometimes better as in our case) accuracies to that of LSTMs. We therefore replaced the LSTM units with GRUs. The GRU unit does not have a forget gate (unlike the LSTM) and has fewer parameters to train on.
- Arguably the most sophisticated classification model, as of now, is the **XLNet** which builds upon the transformer architecture which incorporates encoding and decoding layers in addition to sinusoidal position encoding of words in a sentence. It also incorporates permutations of words to learn more complex relationships between words. It has been pre-trained on a large corpus of data from Wikipedia, BookCorpus, etc. In our experiments, it did improve the validation accuracies.
- To provide an ensemble solution, we combined our best performing models: **XLNet and Bi-GRU with attention**. The final classification decision made by this ensemble was the average of the probabilities coming out of these models. We had lively discussions about whether we should just keep XLNet or do a weighted combination of XLNet with Bi-GRUs. The result was that the argument for robustness in the estimate ensembling these two complex algorithms won the argument of the day. In addition, in our final solution, we kept the Random Forest score as a sanity check output in our display. We believe this model is

- the easiest to explain to lay people. However this score was not used in the average.

Data Science Preprocessing and Modeling

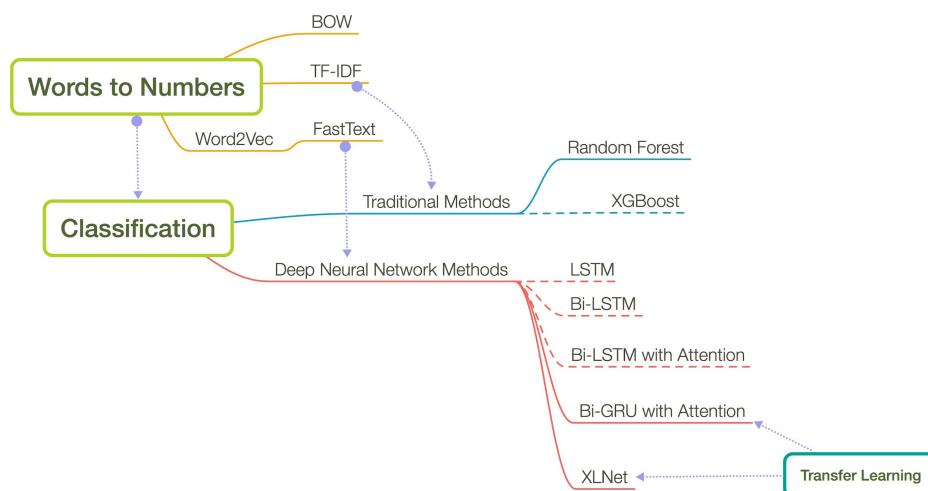


Figure 2: The expanded version of the back-end prediction engine.

CHIEF TECHNOLOGY OFFICER

From the team:

It comprises two major components: Words to Numbers and Classification. Arrows show which words to numbers algorithm used for which classification algorithm. Solid lines represent the models that were integrated into the final solution, while dashed lines represent the ones that we tested during the development process of our final solution.

Thanks!